BREAKOUT GROUP (BOG) 1 REPORT:
TESTBED RESEARCH AGENDAS
AND ADMINISTRATION/USER POLICY ISSUES

1 OBSERVATIONS

1.1 Introduction

Breakout Group (BOG) 1 discussed research agendas and related policy issues for next generation advanced optical network testbeds (ONTs). This BOG 1 report presents a summary of the BOG discussions and the resulting recommendations on ONT research rationales and priorities.

Advanced communication systems and infrastructure constitute a vital national resource. Through the use of advanced optical networking technologies, the potential now exists to create innovative science and application communication environments (i.e., networks, software and systems) that are significantly more powerful than any that exist today.

These new high performance networks and systems will enable the creation of qualitatively and quantitatively new types of communication services, which in turn will provide a basis for the creation of wholly new science paradigms and totally innovative advanced applications. However, to achieve this goal it is necessary to pursue multiyear optical networking research agendas, to provide required research framework processes to support those agendas, and to implement the new generation of ONTs.

Currently, the majority of activity often called “network research” is actually applied engineering oriented toward short-term goals. Even much current real network research is focused on incremental improvements of existing technology and systems. In contrast, the research discussed in this BOG consists of investigations of technologies that are truly innovative and that have the potential for large-scale “disruptive” change, i.e., involving radically fresh innovations as were the Internet and the Web when they were introduced.

The research discussed in this BOG is oriented toward results over longer term time horizons, e.g., 3-5 years rather than 6-12 months. All of the topics are aimed at proof-of-concept and early prototype deployment in government and university ONTs. Several research topics consist of particularly difficult challenges that currently are not being addressed with adequate resources. A few topics represent high-risk investigations, with correspondingly enormous potential for large-scale transformational change. Together, the optical networking research discussed in this BOG represents the next generation of Internetworking, following in the tradition of the government-funded research that led to the Internet and Web.

1.2 The Opportunity

It is not news that important technologies usually require ongoing innovation and development. This observation is especially true for communications services and infrastructure, because these technologies have become the basis for the national economy. What is news is that a set of special circumstances has created a unique window of opportunity for the US to lead the creation of the next wave of Internet innovation by pursuing advanced optical networking research in Federal and partner testbeds.

This window is the threshold of a new era in advanced communications, which occurs only once every fifteen to twenty years. The new era is shaped by three dynamic forcing functions, chief among which are the orders-of-magnitude cost-performance improvements promised by optical network technologies. These technologies will allow us to pursue the second forcing function, i.e., meeting the requirements of a wide array of never-before-achievable large-scale science and engineering (S&E) applications, many of which require multiple terascale data flows in real time among multiple sites worldwide. The third forcing function comes from the new large-scale system architectures that are enabled by optical networking technologies.

The new terascale architectures enabled by these new technologies will enhance or replace much of the existing Internet architecture, with its hierarchical layers of technologies and complex centralized control and management systems. Much of this legacy architecture was based on meeting the requirements of standard voice telephone services—not on the requirements of data services. New architectures will allow for the creation of multiple, differentiated digital communications environments within a common infrastructure. These new architectures will be used to support many types of innovations, such as using distributed communications infrastructure as extended system backplanes, and have the potential of being inherently more secure and robust than the existing Internet.

1.3 The Bad News

Unfortunately, this moment of opportunity has arrived just after a major economic implosion of the telecommunications industry caused by overinvestment in fiber plant which has no visible means of financial support from end-user customers. There are many reasons for this, chief among which are regulatory restrictions on broadband access and the explosive rise of wireless technologies and demand. These reasons notwithstanding, the economic implosion has had a severe negative impact on the normal flow of technologies from industrial research laboratories to commercial products.

Today, industrial research in optical networking communication systems has been almost completely eliminated in favor of applied engineering oriented to near-term product commercialization. Because of the adverse economic conditions, traditional providers have become extremely conservative about acquiring and deploying new technologies. This trend has forced equipment makers to become cautious about new product design and development. Consequently, their device suppliers have either ceased to exist or have focused exclusively on existing product components. This “trickle down” effect of extremely decreased demand is very serious, and has led to a plateauing of US high-performance networking infrastructure.

Although new technologies are still being invented in research labs, these new technologies are frequently orphaned in the laboratories of their birth and early development because the research labs have limited demand for promoting the migration of the new technologies into commercial products. Also, increasingly researchers are asked to address near-term networking issues related to serving existing consumer mass markets, such as the commodity Internet. Due to the confluence of these circumstances, today there are limited opportunities that allow researchers to focus on longer term higher risk or specialized research, such as innovative communication systems in support of advanced science.

These days, the many past contributions of federally funded advanced networking research to US leadership in today’s communication systems, including the Internet, the Web, and voice-over-IP (VoIP) are rarely acknowledged. Although many widely used advanced networking technologies were originally developed for specialized science applications before becoming widely adopted for commercial use, it is much more difficult today to acquire research funding for such projects. If not addressed, these circumstances eventually could become a significant threat to the US economy as a whole, particularly in light of global competition from Europe and Asia.

The remainder of this report focuses on the “good news” that came out of the discussions in this BOG, i.e., the exciting opportunity for US Federal agencies to lead the development of the next Internetworking architecture by aggressively exploiting the existing and planned agency and partner ONTs. The next sections of this report give a summary of the exciting opportunities in this area and the rationale for why US agencies and partner organizations should act immediately, a strawman list of optical network technology research topics that need to be addressed, and the resulting overall recommendation out of this BOG. An appendix contains further observations on process, methods and transformational change.

2 DISCUSSIONS

2.1 Importance of the Research

The "big new promise" of optical internetworking is not in aggregating ever higher levels of existing traffic, but in reducing or eliminating the many existing problems caused by having to share today's network elements. The "aggregation" community is not so much interested in the next milestone in performance as in the growing complexity of the existing Internet, and how to keep the Internet growing in the face of ever increasing traffic volume and complexity demands. On the other hand, future high performance users want to be able to reach deep into the network to control elements of the core, e.g., to obtain a specified guaranteed performance through the core from one laboratory edge network to another. Correspondingly, the future edge networks need to be able to deliver flexible guarantees of performance and specialized services to individual high performance users and systems located in laboratories and data systems connected directly to the edge.

To achieve this, the inherent high performance characteristics of optical fiber need to be combined with the equipment and software needed to bring it all together and make the new services directly available to end users. The main problem to be solved is that, due to the blazing speed of optical networks, making all this work will require researchers to substantially simplify network architectures and protocols rather than adding new layers of complexity. This is the work that needs to be done in the Federal and partner ONTs in the next 3-5 years.

To achieve this performance with simplicity, ONTs are looking into multilayer integrated architectures, featuring direct access to layers and a sophisticated control plane:

- Layer 1: Devices and components
- Layer 2: IEEE 10 Gigabit-per-second Ethernet (Gig-E), 40 Gig-E
- Layer 3: IPv4, IPv6
- Layer 4: Sockets, TCP, UDP

Although drastically simplifying the architecture, nevertheless ONTs need to support both kinds of traffic:

- Aggregation traffic (including support for QoS, priority, preemption . . .)
- Single-stream traffic (e.g., support for mega instruments, supercomputers, terascale databases, haptics, rapid onset of large streams . . .)

2.2 Importance of the Testbeds

Optical network testbeds (ONTs) bring optical networking technology cultures together where they can interact in real-world environments. Technical areas such as optical transport and IP routing/switching are very complex, with great opportunities for synergies to be explored and exploited.

Many synergies are not obvious until complex interactions happen among materials, devices, subsystems, systems, integration, operation, troubleshooting, diagnostics and repair. Many of these interactions eventually lead to new understandings and innovation breakthroughs. ONTs are like Petri dishes for optical internetworking cultures.

Not only government and university researchers, but also startups and larger companies, need "safe haven" testbeds to try out their new ideas in a rigorous but friendly environment. Therefore, in addition to funding government and university research, Federal government agencies should provide support for private sector innovation by encouraging and supporting participation in ONTs by small and large companies.

Startups and smaller companies, by participating in ONTs, can get the early product credibility they need in order to grow to become larger companies, at low or no cost to the government.

Larger communications companies also like to bring pre-release products from their R&D laboratories into safe harbor testbeds, where the R&D people are kept separate from the financial pressures bedeviling the commercial and product people. Large companies realize that they can't sell an idea, and they also understand that validating ideas that originated in their R&D laboratories is not so simple. Ideas that become early product versions may or may not sell in the marketplace. Large companies frequently prefer to bring their pre-release product versions into a noncompetitive environment where the interoperability and functionality kinks can be worked out in cooperation with other players.

Getting pre-release small and large company gear into testbeds frequently costs the government very little beyond establishing the testbed facilities and interconnections and funding the graduate students who provide the bulk of the skilled low-cost labor.

Many future US technologists will earn their masters and PhDs and spend their post-docs working in such testbeds. These long-term investments in people are needed to produce the great companies and teachers for the next generation of US networking competitiveness.

2.3 Why Not Wait for the Carriers?

DARPA/DOE/NSF/NASA funding of university researchers in the 1970s and 1980s allowed for the innovation and incubation that led to the Internet.

In the 1980s, the carriers were late to the Internet party because they understood how to build and sell telephone networks, not the Internet, and the regulatory environment did not understand or appreciate the value of the Internet either. Carrier investments and FCC regulations based on 30-year-amortization schedules for telephone switching equipment were ill-suited for Internet product cycles lasting two or three years at most.

In the 2000s, the carriers will again be late to the ONT party, this time due to massive overbuilding of fiber plant and switch/router infrastructure together with regulatory confusion over build-out for US broadband. This situation has left the carriers focused entirely on where their next quarter's revenue will come from, or whether some of the carriers will even survive into next quarter.

At the same time that the carriers have almost totally ceased to deploy leading-edge optical technologies, multiple localized research efforts within individual research labs, regional optical networks, and advanced optical testbeds continue to drive these technologies forward. However, support for these research efforts has been very limited.

Therefore, it is an appropriate, high-priority role of the US Federal networking lead agencies to establish the conditions in the US to promote the next wave of internetworking innovation by establishing and strongly promoting cooperation and collaboration among the agency- and partner-funded ONTs.

These ONTs will become the technological and architectural foundation of the next Internet, if only the agencies and their partner organizations can understand and act on this exciting situation. Today, the US has a 2-4 year window of opportunity to use the more than 30,000 miles of domestic fiber available to Federal and partner research and education networks at low or no cost to establish the prototype implementations of US optical internetworking.

The role and principal responsibility of the Federal agency networks as well as national R&E networks is to provide their communities with leading-edge production networking services. Therefore, they must ensure that a separate research process exists that continues to evolve new generations of technology within advanced ONTs. This creative process is distinctly different from providing production services engineering, and it cannot be directed by the same types of management techniques. Instead, the research process must be focused on the development of new services and technologies that can then be migrated or morphed into national agency and R&E networks.

2.4 Responsibility of Federal Agency Networks

The lead networking agencies should — must — fund, direct and extend their existing testbeds and programs (a) to leverage the development of the optical networking infrastructure within the US, and (b) to interconnect and collaborate with international ONTs in Europe and Asia in fulfillment of their agency missions:

- **National Science Foundation (NSF)** "big science" (e.g., electronic Very Long Baseline Interferometry—eVLBI) needs much more powerful networks.
- **Department of Defense (DOD)** missions need advanced optical network technologies and architectures (e.g., transformational communications programs).
- **Department of Energy (DOE)** needs to serve its Large Hadron Collider (LHC) and other big science communities (e.g., Scientific Discovery through Advanced Computing—SciDAC —program).
- **National Aeronautics and Space Administration (NASA)** needs to serve its next generation aerospace design and big science communities (e.g., *Columbia* Project supercomputer).
- **National Institutes of Health (NIH)** needs to serve next generation telemedicine and real-time remote medical imaging.

The leading edge of optical networking infrastructure is now moving into a multiyear technology and deployment plateau at the 10-40 Gbps level. This is the perfect level for offering 10 Gbps lambdas (wavelengths) to single-stream and multistream big science applications, e.g., single streams connected to supercomputer backplanes and multistreams connected to 10 Gig-E switches.

The 10 Gbps lambda level will give real scientists using real systems a full two-orders-of-magnitude improvement over today's performance, which for scientists in well-connected laboratories means that today's 15-100 Mbps nominal end-to-end performance will move to the 2.5-10 Gbps level. At such levels of performance, scientists will begin to create **entirely new science**. This happened in the 1980s, when science networks changed from 56 kbps dial-up to 10 Mbps Ethernet, and it happened again in the 1990s with the change from 1.5 Mbps T-1 access connections to 155 Mbps OC-3.

This new science WILL happen again, hopefully in the US, when scientists begin doing real science over ONTs. And when leading end-user scientists in the US see a real change in their ways of doing business, they, as the real users and beneficiaries of the networks, will be the ones to proclaim that the next wave of the Internet has arrived. But this is 5 years away, and only if the US establishes this agenda as a matter of very high priority.

For example, a single 40 Gbps lambda would enable simultaneous support for single-stream 10 Gbps traffic in parallel with aggregation traffic. Parallel 40 Gbps lambdas will enable terascale traffic to be supported; e.g., a 128-node cluster with each node having a 1 Gig-E interface will generate up to 128 Gbps of traffic. How to build and support this type of system is an excellent research topic, and this performance level environment will certainly lead to entirely new science.

When the end-to-end science communities begin to get nominal routine performance in the 1 Gbps and above range, new bottlenecks will then show up in systems and middleware. Research on understanding and then mitigating these bottlenecks is what the ONTs need to focus on in the 3-5 year timeframe. Therefore a systems approach to providing real scientists with **usable** lambda networks is the key to driving the 3-5 year ONT research agendas.

Finally, many agencies, universities and regional optical networks (RONs) are investigating how to assure their access to dark fiber so that they will be able to continue developing and supporting their next-generation R&E networks at this-generation costs. Agencies and partner organizations **should not** plan to build their R&E production networks on today's carrier-dominated and -priced network services. These services are based on legacy technologies and business models that are derived from over a century of focus on telephony services. Also, these infrastructures have all but ceased to evolve. There is a growing gap between the US and many other countries in advanced services provisioning. The existence of this gap may be merely troubling as concerns network services available to US consumers, but this gap is truly alarming as concerns network science and technology research.

The best service provisioning and cost-control strategy for future high-performance R&E networks is to pursue aggressively the optical networking R&D path. The alternative is for agency networks to fall further and further behind in meeting their science and applications networking needs. The bottom line is that agencies and partner organizations need to support ONT testbeds not only as future research vehicles for US competitiveness in the next wave of the Internet but also as the agencies' only affordable growth paths for development of their next-generation production networks.

2.5 Development Cycle of Internet Technologies

In the early days of the Internet in the 1970s and early 1980s, Federal agencies funded university and startup industry researchers to develop the Internet architectures and systems designs, which US industry then developed and deployed in the late 1980s and throughout the 1990s.

The entire development of the Web, beginning with the development of the EU-supported HTML first as a concept and then as a usable tool at CERN in Geneva, followed by the development of the NSF-supported Mosaic freeware at NCSA, followed by Netscape going public and Microsoft distributing its first versions of Internet Explorer, took only a few years. What began as EU and US government-funded research of a few tens of Web servers and a few hundreds or thousands of Web browsers led in less than ten years to the ubiquitous global Web, with "googling" as the information search tool used by millions of people in dozens of languages worldwide.

The important lessons learned from the Internet and the Web are that fresh new internetworking ideas which are considered too radical to have a commercial market emerge from government and partner research laboratories. Then, as the technology capabilities evolve, the ideas and early implementations are noticed and "played around with" by the Internet engineering and user community. The really interesting implementations are quickly picked up and turned into products by entrepreneurial researchers and startup companies. The best of the new ideas and products get the attention of "big" industry, leading to global deployment by companies such as Cisco, Microsoft, and Google, all of which began as visionary startups taking advantage of early implementations of new technologies.

Optical internetworking can and should follow the same development path in the next 3-5 years, with agencies and partner organizations providing the testbed leadership. If agencies do not understand this opportunity and act now, the US could well lose its traditional leadership role in this area. Foreign countries including the EU, Japan, Korea and China – not to mention our good neighbor Canada, which is very competitive in this area – have openly publicized their policies aimed at leapfrogging the US in this area. These countries are funding both the procurement of infrastructure as well as R&D programs seeking to establish themselves as worldwide leaders in optical networking.

US telecom carriers are looking ahead, not several years, but instead only a few quarters at most, because of their desperate competitive posture caused by several years of over-provisioning and regulatory confusion. As a result, over 30,000 miles of fiber-in-the-ground is available to the agency and partner high-performance R&D communities for the next few years at low or no cost. The bottom line is that US agency-funded networking researchers now have a crucial window of opportunity to get the optical internetworking architecture and engineering designs right.

If the US waits for its carriers to develop and deliver optical Internetworking, then just as happened with long-distance telephone service, the carriers will instead dig a hole and bury themselves in today's architectures, because the carriers just do not have the business case to develop a global new paradigm. It is up to government R&D to lead the next wave of the Internet. If the world gets this right, then, just as happened with the earlier waves of the Internet, costs will plummet by several orders of magnitude within just a few years.

The benefits of the next wave of the Internet will extend to everyone, not only the 10-40 Gbps and higher dedicated streams for high-end government big science and applications communities, but also the massive aggregations of 100 Mbps Ethernet for everyone else. There is no better way for the US to get the cost down so well and so quickly than for the US agencies and partner organizations to step up now to lead this area for the next 3-5 years.

2.6 ONT Technology Research Topics

Over the past few years, significant progress has been made in multiple basic research areas related to communications technology. However, because of a set of special circumstances related to the communications marketplace, the natural migration of technology from research labs to commercial deployment has slowed significantly. For some promising technologies, it has completely ceased. Consequently, many of the most promising research innovations remain isolated in labs. This process must be revitalized through optical network testbeds. Below are several areas that should receive enhanced research and development resources in such testbeds.

2.6.1 Layer 4

There are several Layer 4 protocols that can be improved, redesigned, or created anew to provide enhanced support for large-scale data-intensive applications. Preliminary experimentation has demonstrated that integrating these techniques, when implemented with new forms of middleware, can lead to dramatic gains in the types of resource provisioning required by systems designed for next-generation highly distributed infrastructure. Preliminary research has also shown that significant additional advantages can be made possible by integrating new types of L4 protocols with dynamic optical network provisioning.

2.6.2 Layer 3 routing

Remarkable progress has been made on increasing the performance of L3 electronic packet routing. The success of the current widely deployed architecture will always stand as a landmark in the history of communications. Nonetheless, it is clear that, as implemented, this architecture is reaching limits at multiple levels. Models of emerging traffic flow behavior, especially for large-scale flows, clearly indicate that without major changes, required progress will be significantly slowed. L3 traffic today continues to double at a yearly rate. Processor performance doubles at the rate of Moore’s law every 18 months. Major changes are required in numerous L3 areas, including protocols, architecture, processors, interfaces, integration methods, systems, management methods, form factors and power consumption. However, the gains achieved by such changes can be substantially enhanced by integrating these new technologies with emerging optical provisioning methods.

2.6.3 Layer 3 and Layer 2 switching

Similarly, over the past few years, great progress has been made in electronic switching. However, here too, further progress is limited because of multiple architectural and physical factors. Furthermore, it is worth noting that the attempt to completely optimize L3 and L2 functions has been a multiple-year quest that today still has no satisfactory solution. This type of integration is not only a particularly complex challenge, it is also an interdisciplinary challenge, and there have been few opportunities for researchers in these areas to collaborate on solutions. Optical network testbeds would provide major opportunities for such collaboration. Also, such testbeds would assist in promoting some of the most advanced research in this area, that which consists of explorations of the potential for optical components to supplement or replace electronic functionality through new types of hybrid devices.

2.6.4 Wavelength switching

The tremendous potential for advanced optical networking in general, and for wavelength switching in particular, has been widely recognized. Almost all data networking today, including the services provided by national research networks, is supported by legacy technologies such as Synchronous Optical Network (SONET) over static Dense Wavelength Division Multiplexing (DWDM), based on architecture originally designed for voice rather than data communications. Existing advanced optical testbeds are proving the viability of moving beyond this legacy model to much more powerful types of network architecture specifically designed for data communications, e.g., based on more advanced framing methods and dynamic wavelength switching. DWDM can support many hundreds of independent optical channels, each of which can provide 10-40 Gbps transport. The enormous capacity potential of DWDM is well known, but just as important are its capabilities for enhanced control, especially highly distributed control. Advanced systems can be designed so that they can not only individually manage thousands of optical channels, but can also allow the control of those channels to be highly distributed, such as directly distributed to edge processes, even to individual applications. Key research topics for these capabilities include new types of standard device interfaces, such as those being developed by the IETF, not only enhanced Generalized Multiprotocol Label Switching (GMPLS) implementations, but also interfaces that advance beyond signaled overlay architecture.

2.6.5 New edge vs. core paradigms

A new edge vs. core paradigm is fast becoming a key distinction that will fundamentally differentiate tomorrow's networks from today's networks. The current carrier "cloud" model with its centralized control and rigid hierarchical layers, which provides no visibility into core resources, will be replaced by a new model that will provide capabilities for allowing edge processes to discover and dynamically utilize core resources. Today, a major barrier to this type of dynamic utilization of core resources is the current need to use expensive electronic transmitters and receivers to create optical channels so that optical signals can be demultiplexed into components to provide electronic switching. The current generation of passive filters does not adequately address this issue. The eventual solution may consist of all-optical switching, all-optical bypass channels, and all-optical add-drop multiplexors (ADMs). The capabilities made possible by these all-optical technologies will be supplemented by new optical provisioning methods such as dynamic light provisioning and optical burst switching. The investigation of these methods will also assist in the development of optical packet switching technologies. These all-optical components and new methods will significantly enhance restoration and reliability capabilities of networks, including the ability to survive outages even when supporting large scale data flows.

2.6.6 Devices and components

Another wide area of research opportunity consists of investigating many new types of devices, components, and subsystems, such as next generation waveguides, tunable filters, tunable lasers, gate arrays, optical buffers, delay lines, non-linear fiber, and Complementary Metal Oxide Semiconductor (CMOS)-based subsystems. Even as some researchers are investigating how to enhance electronic-based functions by using optical components, other researchers are investigating the addition of electronics to optical components. Optical network testbeds provide opportunities to test such new devices and components in an environment closer to the real world than laboratory experimentation alone.

2.6.7 Management and control planes

These new architectural models will be supported by new types of management and control planes. Almost all communication systems today are based on decades-old, complex, difficult-to-use-and-upgrade legacy systems for management planes, control planes, and Operations Support Systems (OSS). These legacy systems are barriers to new innovations. New research activities on optical network testbeds will enable the development of an entirely new generation of network systems management and control. This research direction is particularly important for enhancing network security.

3 OVERALL BOG 1 RECOMMENDATION

Optical networking is the next Internet technology wave, and therefore, as a national priority, the US should establish the government, university and industry partnerships that will assure US competitiveness in this "new" new area. Therefore, agencies and partner organizations should establish as one of their top priorities during the next three years to educate their leaders, direct their technical teams, and coordinate with their partner organizations, in order to assure US global ONT leadership in the areas of optical networking research, applications including end-to-end issues, and interdomain connection.

4 APPENDIX

4.1 Observations on Process and Methods

- There is a need for new processes related to advanced networking research.
- Advanced networking research must be undertaken within the context of advanced applications; advanced networking research experiments should be integrated with leading-edge applications research.
- Collaborative research should be "grass roots," self-organizing projects among researchers, not top down dictates from administrative hierarchies.

- New means are required for interagency communications on key research topics, driven by the research agenda and the needs of research projects instead of externalities.
- A useful organizing framework for some of these research activities may be Grand Challenge projects.
- Opportunities should exist to undertake high-risk projects.
- Opportunities should exist that encourage interagency experimentation.
- Opportunities should exist for multiple types of advanced testbeds.
- Opportunities should exist for interacting with the global advanced networking research community.
- Opportunities should exist for interdisciplinary research, which has been difficult to implement.

4.2 Observations on Transformational Change

- Research should be oriented toward fundamentally new architectures, not guided by reference to common assumptions and especially not existing legacy systems.
- A fundamentally new model for advanced communications services, systems, and infrastructure is required.
- Research should be oriented to more than increased bandwidth.
- Certainly, high performance is a key requirement, such as multiple terabits per second on individual wavelengths.
- However, many other elements are also important such as flexibility, e.g., capability of mixing traffic exhibiting different behavior characteristics.
- New concepts should be oriented toward a complete systems approach instead of toward individual components.
- New architectures must be designed to enable significantly enhanced flexibility.
- In part, this flexibility should allow for networks capable of dynamic resource utilization and real-time interactive, highly distributed large-scale flows.
- It would also allow for highly distributed control, even at the applications level.
- Research on the design of the core infrastructure should be oriented toward enabling the creation of multiple networks with different properties within distinct environments (e.g., virtual overlay networks).
- Research should be conducted in areas related to ensuring greater level of reliability, capabilities for restoration, automatic self-healing properties, and much higher degrees of security.
- More sophisticated insight into network behavior is required, e.g., “pre” fault problem detection and resolution.
- New types of provisioning models are required.